(1)

Japanese Patent Laid-open Publication No.: HEI11-85412 A

Publication date : March 30, 1999

Applicant : NEC Corporation

Title : DISK ARRAY APPARATUS CONTROLLING METHOD AND DISK ARRAY APPARATUS

(57) [ABSTRACT]

[OBJECT] To prevent the occurrence of erroneous operation in a replacement HDD device and also to detect misconnection, as well, of the HDD device by, when during the operation of a disk array apparatus, a magnetic disk unit (HDD device) has been replaced, constructing so as to automatically detect the status of the HDD device and, simultaneously with this, detect also the defect of a non-volatile memory for storing therein information specifying the defective HDD device, etc.

[SOLVING MEANS] When replacing or repairing a defective HDD device and having re-closed the power source of the HDD device (N_{NEW} numbered HDD device), a disk array controller reads out from the HDD device the characteristic information $A(N_{NEW})$ for discriminating or identifying this HDD device (step 102), reads out from the N_{NEW} numbered non-volatile memory the characteristic information $B(N_{NEW})$ that is stored therein beforehand (step 104), and compares the two with each other (step 106).

[CLAIMS]

[Claim 1] A disk array apparatus controlling method being adapted to control a disk array apparatus having a plurality of HDD devices that store data in the way in which the data is constructed with use of redundancy, the method including steps of:

with respect to a plurality of first non-volatile memories, which respectively

correspond to the HDD devices, storing, for every HDD device, characteristic information used for identifying that HDD device, into the corresponding first non-volatile memory;

storing into a second non-volatile memory, management information used for collectively managing the statuses of the HDD devices;

when during the operation of the disk array apparatus, the HDD device that has failed has been replaced or repaired and the power source of the replacement or repaired HDD device has been re-closed, comparing the characteristic information that is stored in the replacement or repaired HDD device and the characteristic information that is stored in the first non-volatile memory corresponding to the replacement or repaired HDD device and, when those both coincide with each other, starting restoration of that HDD device and, when the both do not coincide with each other, reading out the characteristic information from the replacement or repaired HDD device and writing the characteristic information that has been read out into the first non-volatile memory corresponding to the replacement or repaired HDD device;

when a defect has occurred on any of the HDD devices, when the restoration of the HDD device has been started, and when the restoration of the HDD device has been ended, rewriting the management information within the second non-volatile memory;

when the characteristic information that has been read out from the HDD device is written into the first non-volatile memory, comparing the written content therein and the content that has now been written thereinto and, when those both do not coincide with each other, processing that HDD device as being a defective HDD device; and

when the management information is written into the second non-volatile

memory, comparing the written content there and the content that has now been written there and, when those both do not coincide with each other, processing the disk array apparatus as being a defective disk array apparatus.

[Claim 2] A disk array apparatus controlling method being adapted to control a disk array apparatus having a plurality of HDD devices that store data in the way in which the data is constructed with use of redundancy, the method including steps of:

with respect to a single piece of first non-volatile memory having areas, which respectively correspond to the HDD devices, storing, for every HDD device, characteristic information used for identifying the HDD device;

storing into a second non-volatile memory, management information used for collectively managing the statuses of the HDD devices;

when during the operation of the disk array apparatus, the HDD device that has failed has been replaced or repaired and the power source of the replacement or repaired HDD device has been re-closed, comparing the characteristic information that is stored in the replacement or repaired HDD device and the characteristic information that is stored in the area of the first non-volatile memory corresponding to the replacement or repaired HDD device and, when those both coincide with each other, starting restoration of that HDD device and, when the both do not coincide with each other, reading out the characteristic information from the replacement or repaired HDD device and writing the characteristic information that has been read out into the area of the first non-volatile memory corresponding to the replacement or repaired HDD device and starting the restoration;

when a defect has occurred on any of the HDD devices, when the restoration of the HDD device has been started, and when the restoration of the HDD device has been ended, rewriting the management information within the second non-volatile

memory;

when the characteristic information that has been read out from the HDD device is written into the first non-volatile memory, comparing the written content therein and the content that has now been written thereinto and, when those both do not coincide with each other, processing the disk array apparatus as being a defective disk array apparatus; and

when the management information is written into the second non-volatile memory, comparing the written content there and the content that has now been written there and, when those both do not coincide with each other, processing the disk array apparatus as being a defective disk array apparatus.

[Claim 3] A disk array apparatus controlling method being adapted to control a disk array apparatus having a plurality of HDD devices that store data in the way in which the data is constructed with use of redundancy, a logical drive being set as a collection of the HDD devices constructing respective with-use-of-redundancy data divisions, the method including steps of:

with respect to a plurality of first non-volatile memories, which respectively correspond to the HDD devices, storing, for every HDD device, characteristic information used for identifying that HDD device, into the corresponding first non-volatile memory;

storing into a second non-volatile memory, management information used for collectively managing the statuses of the HDD devices;

when during the operation of the disk array apparatus, the HDD device that has failed has been replaced or repaired and the power source of the replacement or repaired HDD device has been re-closed, comparing the characteristic information that has been read out from the replacement or repaired HDD device and each of all items

of the characteristic information that is stored in the first non-volatile memory and, when no characteristic information that coincides exists, starting restoration of that HDD device and, when there is the characteristic information two items of characteristic information of which coincide with each other, confirming that there do not exist in a respective one of the logical drives, two or more of the HDD devices that are defective and, when there do not exist in a respective one of the logical drives, two or more of the HDD devices that are defective, starting restoration of the replacement or repaired HDD device and, when there exist in any one of the logical drives, two or more of the HDD devices that are defective and when there do not exist two or more of the HDD device that are defective in the logical drive in which there is the first non-volatile memory, the characteristic information of which has coincided with the characteristic information of the replacement or repaired HDD device, starting the restoration of this HDD device and, when, in that case, there exist two or more of the HDD devices that are defective, not performing the restoration of the relevant HDD devices;

when a defect has occurred on any of the HDD devices, when the restoration of the HDD device has been started, and when the restoration of the HDD device has been ended, rewriting the management information within the second non-volatile memory;

when the characteristic information that has been read out from the HDD device is written into the first non-volatile memory, comparing the written content therein and the content that has now been written thereinto and, when those both do not coincide with each other, processing the HDD device as being a defective HDD device; and

when the management information is written into the second non-volatile memory, comparing the written content there and the content that has now been

written there and, when those both do not coincide with each other, processing the disk array apparatus as being a defective disk array apparatus.

[Claim 4] A disk array apparatus controlling method being adapted to control a disk array apparatus having a plurality of HDD devices that store therein data in the way in which the data is constructed with use of redundancy, a logical drive being set as a collection of the HDD devices constructing respective with-use-of-redundancy data divisions, the method including steps of:

with respect to a single piece of first non-volatile memory having areas, which respectively correspond to the HDD devices, storing, for every HDD device, characteristic information used for identifying the HDD device;

storing into a second non-volatile memory, management information used for collectively managing the statuses of the HDD devices;

when during the operation of the disk array apparatus, the HDD device that has failed has been replaced or repaired and the power source of the replacement or repaired HDD device has been re-closed, comparing the characteristic information that has been read out from the replacement or repaired HDD device and each of all items of the characteristic information that is stored in the first non-volatile memory and, when no characteristic information that coincides exists, starting restoration of that HDD device and, when there is the characteristic information two items of characteristic information of which coincide with each other, confirming that there do not exist in a respective one of the logical drives, two or more of the HDD devices that are defective and, when there do exist in a respective one of the logical drives, two or more of the HDD devices that are defective, starting restoration of the replacement or repaired HDD device and, when there exist in any one of the logical drives, two or more of the HDD devices that are defective and when there do not exist two or more of the HDD

device that are defective in the logical drive in which there is the first non-volatile memory, the characteristic information of which has coincided with the characteristic information of the replacement or repaired HDD device, starting the restoration of this HDD and, when, in that case, there exist two or more of the HDD devices that are defective, not performing the restoration of the relevant HDD devices;

when a defect has occurred on any of the HDD devices, when the restoration of the HDD device has been started, and when the restoration of the HDD device has been ended, rewriting the management information within the second non-volatile memory;

when the characteristic information that has been read out from the HDD device is written into the first non-volatile memory, comparing the written content therein and the content that has now been written thereinto and, when those both do not coincide with each other, processing the disk array apparatus as being a defective disk array apparatus; and

when the management information is written into the second non-volatile memory, comparing the written content there and the content that has now been written there and, when those both do not coincide with each other, processing the disk array apparatus as being a defective disk array apparatus.

[Claim 5] A disk array apparatus including:

a plurality of HDD devices that store therein data in the way in which the data is constructed with use of redundancy;

a plurality of first non-volatile memories each of which is provided for every HDD device and stores therein characteristic information used for identifying the corresponding HDD device;

a second non-volatile memory for storing therein management information

used for collectively managing the statuses of the HDD devices;

an HDD device automatic restoration starting unit that, when during the operation of the disk array apparatus, the HDD device that has failed has been replaced or repaired and the power source of the replacement or repaired HDD device has been re-closed, compares the characteristic information that has been read out from the replacement or repaired HDD device and the characteristic information that is stored in the first non-volatile memory corresponding to the replacement or repaired HDD device and, when those both coincide with each other, starts restoration of that HDD device and, when the both do not coincide with each other, reads out the characteristic information from the replacement or repaired HDD device and writes the characteristic information that has been read out into the first non-volatile memory corresponding to the replacement or repaired HDD device, and starts restoration of this HDD device;

a characteristic information write-in unit that reads out the characteristic information from the HDD device and writes it into the first non-volatile memory corresponding to that HDD device;

a management information rewrite unit that, when a defect has occurred on any of the HDD devices, when the restoration of the HDD device has been started, and when the restoration of the HDD device has been ended, rewrites the content of the management information;

a characteristic information write-in defect detection unit that, when the characteristic information that has been read out from the HDD device is written into the first non-volatile memory, compares the written content therein and the content that has now been written thereinto and, when those both do not coincide with each other, processes that HDD device as being a defective HDD device; and

a management information write-in defect detection unit that, when the management information is written into the second non-volatile memory, compares the written content there and the content that has now been written there and, when those both do not coincide with each other, processes the disk array apparatus as being a defective disk array apparatus.

[Claim 6] A disk array apparatus including:

a plurality of HDD devices that store therein data in the way in which the data is constructed with use of redundancy;

a single piece of first non-volatile memory that, for storing therein characteristic information used for identifying the HDD devices, is commonly provided with respect to the HDD devices and that has areas each corresponding to a respective one of the HDD devices;

a second non-volatile memory for storing therein management information used for collectively managing the statuses of the HDD devices;

an HDD device automatic restoration starting unit that, when during the operation of the disk array apparatus, the HDD device that has failed has been replaced or repaired and the power source of the replacement or repaired HDD device has been re-closed, compares the characteristic information that is stored in the replacement or repaired HDD device and the characteristic information that is stored in the area corresponding to the relevant HDD device in the first non-volatile memory and, when those both coincide with each other, starts restoration of that HDD device and, when the both do not coincide with each other, reads out the characteristic information from the replacement or repaired HDD device and writes the characteristic information that has been read out into the area corresponding to the replacement or repaired HDD device, of the first non-volatile memory, and starts restoration of this HDD device;

a characteristic information write-in unit that reads out the characteristic information from the HDD device and writes it into the area corresponding to this HDD device, of the first non-volatile memory;

a management information rewrite unit that, when a defect has occurred on any of the HDD devices, when the restoration of the HDD device has been started, and when the restoration of the HDD device has been ended, rewrites the content of the management information;

a non-volatile memory information write-in defect detection unit that, when the characteristic information that has been read out from the HDD device is written into the first non-volatile memory, compares the written content therein and the content that has now been written thereinto and, when those both do not coincide with each other, processes the disk array apparatus as being a defective disk array apparatus; and

a management information write-in defect detection unit that, when the management information is written into the second non-volatile memory, compares the written content there and the content that has now been written there and, when those both do not coincide with each other, processes the disk array apparatus as being a defective disk array apparatus.

[Claim 7] A disk array apparatus including:

a plurality of HDD devices that store therein data in the way in which the data is constructed with use of redundancy;

a plurality of first non-volatile memories each of which is provided for every HDD device and stores therein characteristic information used for identifying the corresponding HDD device;

a second non-volatile memory for storing therein management information used for collectively managing the statuses of the HDD devices;

a replacement HDD device automatic restoration start determination unit that, when during the operation of the disk array apparatus, the HDD device that has failed has been replaced or repaired and the power source of the replacement or repaired HDD device has been re-closed, compares the characteristic information that has been read out from the replacement or repaired HDD device and each of all the characteristic information that is stored in the first non-volatile memory and, when no characteristic information that coincides exists, starts restoration of that HDD device and, when there is the characteristic information two items of characteristic information of which coincide with each other, confirms that there do not exist in a respective one of the logical drives, two or more of the HDD devices that are defective and, when there do not exist in a respective one of the logical drives, two or more of the HDD devices that are defective, starts restoration of the replacement or repaired HDD device and, when there exist in any one of the logical drives, two or more of the HDD devices that are defective and when there do not exist two or more of the HDD device that are defective in the logical drive in which there is the first non-volatile memory, the characteristic information of which has coincided with the characteristic information of the replacement or repaired HDD device, starts the restoration of this HDD and, when, in that case, there exist two or more of the HDD devices that are defective, does not perform the restoration of the relevant HDD devices;

a characteristic information write-in unit that reads out the characteristic information from the HDD device and writes it into the first non-volatile memory corresponding to that HDD device;

a management information rewrite unit that, when a defect has occurred on any of the HDD devices, when the restoration of the HDD device has been started, and when the restoration of the HDD device has been ended, rewrites the content of

the management information;

a characteristic information write-in defect detection unit that, when the characteristic information that has been read out from the HDD device is written into the first non-volatile memory, compares the written content therein and the content that has now been written thereinto and, when those both do not coincide with each other, processes that HDD device as being a defective HDD device; and

a management information write-in defect detection unit that, when the management information is written into the second non-volatile memory, compares the written content there and the content that has now been written there and, when those both do not coincide with each other, processes the disk array apparatus as being a defective disk array apparatus, the logical drive being set as a collection of the HDD devices constructing respective with-use-of-redundancy data divisions.

[Claim 8] A disk array apparatus including:

a plurality of HDD devices that store therein data in the way in which the data is constructed with use of redundancy;

a single piece of first non-volatile memory that, for storing therein characteristic information used for identifying the HDD devices, is commonly provided with respect to the HDD devices and that has areas each corresponding to a respective one of the HDD devices;

a second non-volatile memory for storing therein management information used for collectively managing the statuses of the HDD devices;

a replacement HDD device automatic restoration start determination unit that, when during the operation of the disk array apparatus, the HDD device that has failed has been replaced or repaired and the power source of the replacement or repaired HDD device has been re-closed, compares the characteristic information that is stored

in the replacement or repaired HDD device and each of all the characteristic information that is stored in the first non-volatile memory and, when no characteristic information that coincides exists, starts restoration of that HDD device and, when there is the characteristic information, two items of characteristic information of which coincide with each other, confirms that there do not exist in a respective one of the logical drives, two or more of the HDD devices that are defective and, when there do not exist in a respective one of the logical drives, two or more of the HDD devices that are defective, starts restoration of the replacement or repaired HDD device and, when there exist in any one of the logical drives, two or more of the HDD devices that are defective and when there do not exist two or more of the HDD device that are defective in the logical drive in which there is the first non-volatile memory, the characteristic information of which has coincided with the characteristic information of the replacement or repaired HDD device, starts the restoration of this HDD and, when, in that case, there exist two or more of the HDD devices that are defective, does not perform the restoration of the relevant HDD devices;

a characteristic information write-in unit that reads out the characteristic information from the HDD device and writes it into the area in the first non-volatile memory corresponding to that HDD device;

a management information rewrite unit that, when a defect has occurred on any of the HDD devices, when the restoration of the HDD device has been started, and when the restoration of the HDD device has been ended, rewrites the content of the management information;

a non-volatile memory information write-in defect detection unit that, when the characteristic information that has been read out from the HDD device is written into the first non-volatile memory, compares the written content therein and the content that

has now been written thereinto and, when those both do not coincide with each other, processes the disk array apparatus as being a defective disk array apparatus; and

a management information write-in defect detection unit that, when the management information is written into the second non-volatile memory, compares the written content there and the content that has now been written there and, when those both do not coincide with each other, processes the disk array apparatus as being a defective disk array apparatus, the logical drive being set as a collection of the HDD devices constructing respective with-use-of-redundancy data divisions.

[0021]

[PROBLEMS THAT THE INVENTION IS TO SOLVE] However, in the above-described conventional disk array system, when having replaced a defective HDD device, there is adopted a method of once turning off the power source of the disk array apparatus, then rebooting the disk array apparatus as a whole, and thereby starting restoration of that relevant HDD device. However, most of the disk array apparatuses each have a plurality of logical drives and there are cases where those apparatuses are being operated over a period of 24 hours. In those cases, in the case of each of the above-described disk array systems, it happens that, even when a defective HDD device occurs in one logical drive, the power source of the disk array apparatus cannot be turned off. Resultantly, there occurs the inconvenience that the restoration of the defective HDD device cannot be performed.

[0022] In addition, when defective HDD devices exist within a plurality of logical drives within one disk array apparatus and when an attempt is made to replace those defective HDD devices, it is considered as being possible that other than the defective HDD devices, a normal HDD device may be drawn off. Since the HDD device that

has been erroneously drawn away normally operates in repair or in examination, it is attempted to return it to the original position of the disk array apparatus. Although, at this time, no problem occurs where another defective HDD device doesn't exist within the logical drive from which the normal HDD device that has been erroneously drawn away exists, when such a defective HDD device exists, it results that the with-use-of-redundancy data construction is lost. In this case, unless the data of the HDD device that has been erroneously drawn off is changed, restoration of the data can be restored through the use of manpower. However, when having loaded that normal HDD device with respect to a position of loading of another logical drive, since that HDD device normally operates, it is possible that one will start the restoration as the system is. However, in this case, the with-use-of-redundancy data construction of the logical drive in which the HDD devices that are all normal that has until then existed collapses. This leads to the inconvenience that even restoration using manpower becomes impossible.

[0023] It is the object of the present invention to improve the above-described inconvenience that the conventional disk array apparatus has, especially to provide a disk array apparatus that, during the operation of the disk array apparatus, enables appropriately executing the restoration processing according to the status of the HDD devices that are relevant thereto.

[0024] Concretely, when during the operation of the disk array apparatus, an HDD device has been replaced, there is adopted a method of automatically detecting the status of the HDD device and, simultaneously with this, also detecting defects of the non-volatile memory, as well, thereby preventing the occurrence of an erroneous operation of the replacement HDD device leading to enhancement of the reliability of the system. Also, by detecting the misconnection, as well, of the HDD devices, it is

aimed to prevent the relevant data from being lost due to the human mistakes, thereby enhancing the reliability of the system.

[0038]

[Embodiment of the Invention] Next, a preferred embodiment of the present invention will be explained with reference to the drawings. Fig. 1 is a block diagram that shows the construction of a disk array apparatus according to an embodiment of the present invention.

[0039] A disk array apparatus 11 shown in Fig. 1 is equipped with an interface 12 that is used for connection thereof with a higher-order device 21, a disk array controller 13, an n number of HDD devices 14_1 to 14_n, non-volatile memories 15_1 to 15_n that correspond to the respective HDD devices 14_1 to 14_n, and a non-volatile memory 19 that stores therein management information. A request to access (a request to read out or a request to write in) the disk array apparatus 11 from the higher-order device 21 is input to the disk array controller 13 via the interface 12.

[0040] The disk array controller 13 has a mechanism that, according to the content of the request from the higher-order device 21, controls the respective HDD devices 14_1 to 14_n and performs read-out or write-in of data, a mechanism that, according to the information within the non-volatile memory 19, determines the status of the HDD devices 14_1 to 14_n, and a mechanism that, according to the information within the non-volatile memories 15_1 to 15_n and the characteristic information 20_1 to 20_n of the HDD devices 14_1 to 14_n, determines the status of the HDD device. Incidentally, these mechanisms of the disk array controller 13 are constructed of a processor for use for performing control, a program memory having stored therein a program that defines that controlling processor, a memory having stored therein parameters that are used

when the apparatus operates, an interface for performing read-out or write-in with respect to the HDD devices 14_1 to 14_n, an interface that is used with respect to the non-volatile memories 15_1 to 15_n corresponding to the HDD devices 14_1 to 14_n respectively, and an interface used with respect to the non-volatile memory 19 for storing therein the management information. The characteristic information 20_1 to 20_n are respectively stored within the HDD devices 14_1 to 14_n.

[0041] This disk array apparatus 11 is a type with respect to which one or a plurality of logical drives can be set. The logical drive is the one a single piece of which is set as a collection of the HDD devices that construct their respective with-use-of-redundancy data divisions.

[0042] First, using a flow chart of Fig. 2, the operation of this disk array apparatus the logical drive of which is one in number will be explained, taking up as an example a process that is executed when during the operation of the disk array apparatus, the power source of a single piece of HDD device has been re-closed.

[0043] In this disk array apparatus, for managing the respective HDD devices, non-volatile memories, etc., parameters such as a counter i, and data arrangements A(i), B(i), C(i), and D(i) are used. The counter i is a parameter for identifying the HDD device. A characteristic information arrangement that stores characteristic information of the HDD devices, characteristic information arrangement that stores the characteristic information in the non-volatile memories, management information arrangement that stores management information for a defective one of the HDD devices, and management information arrangement that stores management information for one of the HDD devices that is destined for restoration are represented by A(i), B(i), C(i), and D(i), respectively. For temporarily storing these parameters there is used a memory within the disk array controller 13. Also, in the disk array

controller 13, an HDD device automatic restoration starting unit that starts automatic restoration of the HDD devices 14_1 to 14_n, a characteristic information write-in unit that reads out characteristic information from the HDD devices 14_1 to 14_n and writes it into the non-volatile memories 15_1 to 15_n, a management information rewrite unit that rewrites the management information stored in the non-volatile memory 19, a characteristic information write-in defect detection unit that detects a defect at the time of white of characteristic information into the non-volatile memories 15_1 to 15_n, and a management information write-in defect detection unit that detects a defect at the time of write of management information into the non-volatile memory 19, are constructed by the above-described controlling processor.

[0044] When during the operation of the disk array apparatus, the power source of the HDD device has been re-closed, the disk array controller 13, first, sets the identification number of the HDD device, the power source of which has been re-closed, to be the parameter N_{NEW} and sets "1" to the counter i (step 101). And, the disk array controller 13 reads out the characteristic information A (N_{NEW}) from the N_{NEW} numbered HDD device (the HDD device the power source of which has been re-closed) (step 102), and determines whether that read-out has normally been performed (step 103). When the read-out of the characteristic information A (N_{NEW}) from the N_{NEW} numbered HDD device has not normally been performed, the operation flow proceeds to step 112. When having normally read out the characteristic information A (N_{NEW}), the controller 13 reads out the characteristic information B (N_{NEW}) from the N_{NEW} numbered non-volatile memory (a relevant one of the non-volatile memories 15_1 to 15_n) (step 104) and determines whether the read-out from the non-volatile memory 15 has normally been performed (step 105).

[0045] When at step 105 the characteristic information B(N_{NEW}) has not normally been

read out from the N_{NEW} numbered non-volatile memory, the flow proceeds to step 112. On the other hand, when the characteristic information $B(N_{NEW})$ has normally been read out from the N_{NEW} numbered non-volatile memory, the controller 13 compares the characteristic information $A(N_{NEW})$ of the N_{NEW} numbered HDD device and the characteristic information $B(N_{NEW})$ of the N_{NEW} numbered non-volatile memory (step 106). When the both coincide, the flow proceeds to step 109. When at step 106 the $A(N_{NEW})$ and the $B(N_{NEW})$ do not coincide with each other, the characteristic information $A(N_{NEW})$ of the N_{NEW} numbered HDD device is set as the characteristic information $B(N_{NEW})$ of the N_{NEW} numbered non-volatile memory and the characteristic information $B(N_{NEW})$ is written into the N_{NEW} numbered non-volatile memory (step 107) and it is determined whether this write-in has normally been performed (step 108). When the write-in of the characteristic information $B(N_{NEW})$ into the N_{NEW} numbered non-volatile memory has normally been performed, the flow proceeds to step 109. When that write-in has not normally been performed, the flow proceeds to step 112. When that write-in has not normally been performed, it means that the status of the N_{NEW} numbered HDD device is the one that has existed before the re-closure of the power source thereof, i.e., the status of the defective HDD device.

[0046] At step 109, "1" is set as the management information $D(N_{NEW})$ indicating the status of restoration of the N_{NEW} numbered HDD device destined for restoration, and that management information $D(N_{NEW})$ is written into the non-volatile memory 19 that stores the management information, and it is determined whether that write-in has normally been performed (step 110). When the management information $D(N_{NEW})$ indicating the status of restoration has not normally been written into the non-volatile memory 19 storing the management information therein, the operation flow proceeds to step 112. On the other hand, when that write-in has normally been performed at

step 110, the operation flow proceeds to step 111, in which the restoration is started.

[0047] At step 112, the controller 13 reports an error message and terminates the process. At this time, when the characteristic information $A(N_{NEW})$ has failed to be read out from the N_{NEW} numbered HDD device, and read-out and write-in with respect to the N_{NEW} numbered non-volatile memory have failed to be performed, they correspond to the status of the N_{NEW} numbered HDD device that has existed before re-closure of the power source, i.e. the status where the HDD device is defective. When write of the management information $D(N_{NEW})$ into the non-volatile memory 19 has failed to be performed, it is defined that the disk array apparatus be defective. The method of reporting the error message includes a method of reporting the status to the higher-order device 21 and a method of making an LED display. However, any method is more preferable if that method can make it acknowledged which position of the defective HDD device is located in or that the disk array apparatus is defective.

[0048] In this disk array apparatus, as the management information C(i) indicating the status of the HDD device that is defective, written into the non-volatile memory 19, and the management information D(i) indicating the status of the HDD that is under restoration, a bit indicating that the HDD device is defective and a bit indicating that the HDD device is under restoration are used, respectively, in corresponding relationship to those HDD devices. When all the HDD devices are normally operating, the bit representing the management information C(i) indicating the defective HDD device and the bit representing the management information D(i) indicating the HDD device under restoration are all "0". When there is a defective HDD device, the bit representing the management information C(i) that corresponds to that defective HDD device is "1" and the bits representing the management information D(i) indicating the HDD device under restoration are all "0". Also, regarding the HDD device that is

under restoration, although the bits representing the management information D(i) indicating the HDD device under restoration are all "1", the HDD devices that are under restoration and the bits corresponding to that are "1" are only one in number.

[0049] Also, in this disk array apparatus, as the characteristic information A(i) of the HDD device and the characteristic information B(i) of the non-volatile memory 15, there is used 36-byte inquiry data that includes the vendor name and serial number. However, the data representing the characteristic information is not limited thereto. For example, using a random number, using a slot number, or using information that represents date and time is possible. However, in such cases, it is preferable to use data representing the characteristic information corresponding to the HDD device, the use of which reliably discriminates the characteristic information for every HDD device.

[0050] Next, using flow charts of Figs. 3 and 4, the operation of this disk array apparatus when the apparatus has two or more logical drives will be explained, taking up as an example a process that is executed when during the operation of the disk array apparatus, the power source of one HDD device has been re-closed.

[0051] Figs. 3 and 4 as a whole show a single sheet of flow chart and, by the encircled symbols P and Q, the mutual connection relationship between the two is shown, and those figures show the flow of the operation of the disk array controller 13.

[0052] Here, for managing the respective HDD devices, non-volatile memories, etc., there are used parameters such as counter i, counter j, counter N_J, and counter N_{EQUAL} and data arrangements A(i), B(i), C(i), D(i) and E(i). The counter i is a parameter for identifying an HDD device and the counter j is a parameter for identifying a logical drive within the apparatus. The counter j is set in the way of being related to the counter i used for identifying the HDD device, depending on the conditions on which the apparatus is constructed. A characteristic information arrangement that stores

characteristic information of the HDD devices, characteristic information arrangement that stores the characteristic information in the non-volatile memories, management information arrangement that stores management information for a defective one of the HDD devices, management information arrangement that stores management information for one of the HDD devices that is destined for restoration, and arrangement E(i) used for identifying the number of the HDD device whose power source has been re-closed and whose characteristic information has coincided with that stored in the non-volatile memory are represented by A(i), B(i), C(i), D(i), and E(i), respectively. Further, the counter used for investigating the number of the defective HDD devices located within the j-numbered logical drive is represented by N_J, and the temporary counter is represented by N_{EQUAL}. For temporal storage of those parameters, a memory within the disk array controller 13 is used. Also, in the disk array controller 13, a replacement HDD device restoration starting unit that starts automatic restoration of the HDD devices 14_1 to 14_n, a characteristic information write-in unit that reads out characteristic information from the HDD devices 14_1 to 14_n and writes it into the non-volatile memories 15_1 to 15_n, a management information rewrite unit that rewrites the management information stored in the non-volatile memory 19, a characteristic information write-in defect detection that that detects a defect at the time of write of the characteristic information into the non-volatile memories 15_1 to 15_n, and a management information write-in defect detection unit that detects a defect at the time of write of management information into the non-volatile memory 19, are constructed by the above-described controlling processor.

[0053] When during the operation of the disk array apparatus, the power source of the HDD device has been re-closed, the disk array controller 13, first, sets the identification number of the HDD device, the power source of which has been re-closed, to be the

parameter N_{NEW} and sets "1" to the HDD device identification counter i. With respect to the arrangement E(i) used for identifying the number of the HDD device whose power source has been re-closed and whose characteristic information has coincided with that stored in the non-volatile memory, "0" is set , and "0" is set with respect to the counter N_J used for investigating the number of the defective HDD devices located within the j-numbered logical drive, and "0" is set with respect to the temporary counter N_{EQUAL} (step 201). And, the disk array controller reads out from the N_{NEW} numbered HDD device the characteristic information $A(N_{NEW})$ that corresponds thereto (step 202) and determines whether that read-out has normally been performed (step 203). Here, when that read-out has not normally been performed, the controller 13 reports to the higher-order device 21 an error message to the effect that the relevant HDD device is a defective one (step 317) and terminates the processing. When the read-out of the characteristic information $A(N_{NEW})$ from the N_{NEW} numbered HDD device has normally been performed, the controller 13 reads out the characteristic information B(i) from the i-numbered non-volatile memory 15 (step 204) and determines whether read-out from that non-volatile memory 15 has normally been performed (step 205).

[0054] When at step 205 the read-out of the characteristic information B(i) from the i-numbered non-volatile memory 15 has not normally been performed, the operation flow proceeds to step 210. On the other hand, when that read-out has normally been performed at step 205, the controller 13 compares the characteristic information $A(N_{NEW})$ of the HDD device, the power source of which has been re-closed and the characteristic information B(i) that is stored in the i-numbered non-volatile memory (step 206). When those both coincide, "1" is set to the E(i), "1" is added to the temporary counter N_{EQUAL}, and "1" is added to the counter N_J for investigating the number of the defective HDD devices located within the j-numbered logical drive (step

216), then the flow proceeds to step 217. In contrast to this, when at step 206 the $A(N_{NEW})$ and B(i) characteristic information do not coincide, the controller 13 reads out the management information C(i) for the defective HDD device from the non-volatile memory 19 (step 207) and determines whether that read-out has normally been performed. When the read-out of the management information C(i) for the defective HDD device from the non-volatile memory 19 has not normally been performed, the controller 13 reports to the higher-order device 12 an error message to the effect that the current disk array apparatus is a defective one (step 317), and terminates this processing. When at step 208 the read-out of the management information C(i) has normally been performed, the controller 13 determines whether that management information C(i) indicating a defective HDD device is "1" (step 209). If the C(i) = "1", the flow proceeds to step 215. If at step 209 C(i) ≠ "1", the flow proceeds to step 217.

[0055] At step 210, the controller 13 reads out the management information C(i) for the defective HDD device from the non-volatile memory 19 and determines at step 211 whether that read-out has normally been performed. Here, when read-out of the management information C(i) from the non-volatile memory 19 has not normally been performed, the controller 13 reports to the higher-order device 21 an error message to the effect that the current disk array apparatus is a defective one (step 317), and terminates the processing. On the other hand, when at step 211that read-out has normally been performed, the controller 13 determines whether that management information C(i) indicating a defective HDD device is "1" (step 212). If the C(i) = "1", the flow proceeds to step 215. If C(i) ≠ "1", "1" is set as the management information C(i) indicating a defective HDD device, the controller 13 writes that C(i) into the non-volatile memory 19 (step 213), and determines whether that write-in has normally been performed (step 214). When the write-in of that management information C(i)

has normally been performed, the flow proceeds to step 215. When having not normally been performed, the controller 13 reports to the higher-order device 21 an error message to the effect of being a defective disk array apparatus (step 317), and terminates this processing.

[0056] At step 215, "1" is added to the counter N_J for investigating the number of defective HDD devices existing within the j-numbered logical drive, then the flow proceeds to step 217. The logical drive identification counter j is derived from the apparatus construction conditions and HDD device number i.

[0057] At step 217, in order to repeat the above-described preceding processing operations with respect to a respective one of all the HDD devices, "1" is added to the counter i for identifying the HDD devices. And the disk array controller 13 determines whether the value of the counter i after that addition has been performed is equal to or smaller than the number of all the HDD devices (step 218). If the value of the i is equal to or smaller than the number of all the HDD devices, the flow proceeds to step 204. If the value of the counter i is greater than the number of the HDD devices, the flow proceeds to step 301.

[0058] At step 301, the controller 13 determines whether the temporary counter N_{EQUAL} is "0". Here, if the N_{EQUAL}="0", the flow proceeds to step 312. In this case, it results that the characteristic information $A(N_{NEW})$ of the HDD device the power source of which has been re-closed is different from any one item of information of the characteristic information B(i) stored in the non-volatile memory. On the other hand, if at step 301 $N_{EQUAL} \neq$ "0", "1" is set to the counter j for identifying the logical drive within the disk array apparatus and "0" is set to the temporary counter N_{EQUAL} (step 302). Then, the number of the first HDD device of the logical drive j is set with respect to the counter i for identifying the HDD (step 303), then the flow proceeds to step 304.

[0059] At step 304, the disk array controller 13 determines whether the arrangement E(i) used for identifying the number of the HDD device whose power source has been re-closed and whose characteristic information has coincided with that stored in the non-volatile memory is "1". If the $E(i) \neq 1$, "1" is added to the counter i for identifying the HDD device (step 305), then the controller 13 determines whether the counter i is not above the final HDD device of the logical drive j (step 306). If the counter i is not above the final HDD device of the logical drive j, in order to repeat the above-described preceding processing operations with respect to a respective one of the HDD devices that belong to the logical drive j, the flow returns to step 304. If the counter i is above the final HDD device of the logical drive j, the flow proceeds to step 309. In addition, if at step 304 the E(i) is "1", it is determined whether the number N_J of the defective HDD devices existing within the j-numbered logical drive is more than two (step 307). If the N_J is two or more, after "1" is added to the temporary counter N_{EQUAL} (step 308), a transfer is made to step 309. If the N_J is less than two, the flow proceeds to step 309 intact.

[0060] At step 309, "1" is added to the counter j for identifying the logical drive, then it is determined at step 310 whether the value of the j is equal to or smaller than the number of all the logical drives. If the j is equal to smaller than the number of all the logical drives, the disk array controller 13 repeats the processing operations that are performed from step 303. If the j is above the number of all the logical drives, it is checked whether the temporary counter N_{EQUAL} is "0" (step 311). When the N_{EQUAL} is not "0", it falls upon either a case where two or more defective HDD devices exist within the logical drive having the N_{NEW} numbered HDD device located therewithin and therefore restoration is impossible or a case where, other than the logical drive within which the N_{NEW} numbered HDD device exists, there is a defective HDD device the

characteristic information of which is the same as that of the N_{NEW} numbered HDD device and the number of which is two or more and therefore it is likely that the with-use-of-redundancy data construction will collapse. Therefore, the controller 13 reports an error message to the higher-order device 21 (step 317) and terminates this processing. If at step 311 the temporary counter N_{EQUAL} is "0", the flow proceeds to step 312.

[0061] At step 312, the characteristic information $A(N_{NEW})$ of the N_{NEW} numbered HDD device is set with respect to the characteristic information $B(N_{NEW})$ of the NNEW numbered non-volatile memory, and the characteristic information $B(N_{NEW})$ of the N_{NEW} numbered non-volatile memory is written thereinto. And at step 313, it is determined whether write of the characteristic information $B(N_{NEW})$ into the N_{NEW} numbered non-volatile memory has normally been performed. When that write-in has normally been performed, the flow proceeds to step 314. When that write-in has not normally been performed, the controller 13 reports an error message to the higher-order device 21 (step 317), and then terminates the processing.

[0062] At step 314, "1" is set with respect to the management information $D(N_{NEW})$ that represents the status of being restored of the N_{NEW} numbered HDD device destined for restoration, thereby that management information is written into the non-volatile memory 19 that stores therein the management information. And it is determined whether write of the management information $D(N_{NEW})$ that represents the status of being restored into the non-volatile memory 19 that stores therein the management information has normally been performed (step 315). When that write has not normally been performed, the controller 13 reports to the higher-order device 21 an error message to the effect of being a defective disk array apparatus (step 317) and then terminates the processing. On the other hand, when having normally been

performed at step 315, the controller 13 starts restoration of the HDD device (step 316), and then terminates the processing operations.

[0063] The disk array apparatus according to this embodiment that has been explained above has a construction of being equipped with the non-volatile memories for storing therein the characteristic information correspondingly to the individual HDD devices and of using, separately from those memories, the non-volatile memory for storing therein the management information. However, the construction of the disk array apparatus to which the present invention is applied is not limited to that construction. For example, it may be arranged that the non-volatile memories for storing therein the characteristic information be unified into one non-volatile memory, thereby the storing area of this non-volatile memory is divided into regions for storing therein the characteristic information of the individual HDD devices. Or it may be arranged that the management information and characteristic information be all stored into a single piece of non-volatile memory in the way in which this memory area is divided into storage regions corresponding to those items of information. When it has been arranged that all the characteristic information be stored into a single piece of non-volatile memory, the error message that is to be reported to the higher-order device is used as the one the error of which means a defective disk array apparatus. The reason for this is as follows. When the non-volatile memory for storing the characteristic information has been made one in number, failing to perform read-out or write-in with respect to one region within that non-volatile memory results in that the operation of that non-volatile memory itself is abnormal. Therefore, even if succeeding in performing read-out or write-in with respect to the other regions of that memory, the reliability of such data becomes seriously decreased. Therefore, it is better not to use that memory. Concretely, in the disk array controller 13, it is arranged

that, instead of the characteristic information write-in defect detection unit that processes the error message as being a defective HDD device, there be provided a non-volatile memory information write-in defect detection unit that processes the error message as being a defective disk array apparatus.

[0064] Also, in the disk array apparatus of the above-described embodiment, it adopts a with-use-of-redundancy data construction enabling the restoration of only one HDD device. However, it is possible to increase the parameters for discriminating defective HDD devices in corresponding relationship to the number of the defective HDD devices that can be restored and to execute control that is performed, according to the branch conditions, using the counter that indicates the number of the defective HDD devices. By doing so, it is possible to adopt a with-use-of-redundancy data construction that enables restoration of two or more defective HDD devices. In addition, in the above-described disk array apparatus, it is arranged that, when re-closing the power source of a defective HDD device, restoration be performed with respect to the part of the HDD devices which can automatically be restored. However, without providing that mechanism, it is possible to construct so that the restoring operation may be started in accordance with the instruction command that is issued from the higher-order device.

[0065]

[Effect of the Invention] As has been described above, according to the present invention, when an HDD device has been replaced during the operation of the disk array apparatus, the status of that replacement HDD device is automatically detected and, simultaneously therewith, the defect of the non-volatile memory also is detected. As a result of this, the effect is brought about that the erroneous operation of the replacement HDD device can be prevented from occurring. In addition, by detecting

the misconnection, as well, of the HDD device, the effect is brought about that relevant data can be prevented from being lost due to the human mistakes.